DOCUMENT RESUME

ED 042 138 AL 002 441

AUTHOR	Kay , Martin
TITLE	Performance Grammars.
REPORT NO	P-4391
PUB DATE	Jun 70
NOTE	18p.; Paper presented at the National Science Foundation Seminar on the Construction of Complex Grammars, Harvard University, Cambridge, Mass., June 1970
AVAILABLE FROM	Reports Dept., The Rand Corporation, 1700 Main Street, Santa Monica, Calif. 90406 ($1.00)
EDRS PRICE	EDRS Price MF-$0.25 HC-$1.00
DESCRIPTORS	*Computational Linguistics, Grammar, Language Universals, Linguistic Competence, *Linguistic Performance, *Linguistic Theory, Mathematical Linguistics, Models, *Semantics, Sentences, Surface Structure

ABSTRACT

The author outlines the construction of a somewhat different machine than that envisioned by Turing (with which it would be possible to converse, presumably by telephone or telegraph, and which would be capable of masquerading as a human being). The machine envisaged by the author would be capable of doing comprehension exercises. Such a machine that is to understand what people say must be capable of analyzing the grammatical structures of the utterances it receives, presenting the information that they contain in some kind of canonical form, of storing expressions in the canonical form in such a way that they can easily be retrieved, of locating those expressions that are relevant to a given question, and of constructing answers that accord with the rules of English grammar. If it is in fact possible to make the machine find the deep structures that correspond to the surface strings it receives, then the efficiency of all subsequent operations will be immeasurably increased, perhaps enough to make practically feasible a project that otherwise would have been totally unthinkable. (AMM)

ED042138

U.S. DEPARTMENT OF HEALTH, EDUCATION & WELFARE
OFFICE OF EDUCATION
THIS DOCUMENT HAS BEEN REPRODUCED EXACTLY AS RECEIVED FROM THE PERSON OR ORGANIZATION ORIGINATING IT. POINTS OF VIEW OR OPINIONS STATED DO NOT NECESSARILY REPRESENT OFFICIAL OFFICE OF EDUCATION POSITION OR POLICY.

PERFORMANCE GRAMMARS

Martin Kay

June 1970

AL 002 441

P-4391

PERFORMANCE GRAMMARS

Martin Kay*

The RAND Corporation, Santa Monica, California

The study of linguistic competence, as applied to a particular language, is the attempt to characterize that subset of the Cartesian product of all possible meanings with all possible utterable sequences of sounds that are the correct sentences of the language. By talking of the correct sentences we do not, of course, mean to suggest that the study of linguistic competence is, in any way, involved with social norms or with rhetoric; only that it is concerned with the pairs of sounds and meanings that speakers recognize as belonging to their language rather than with the often imperfect imitations of these that human imperfection forces them to be content with in practice. The characterization of the meaning–sound pairs that constitute the language is particularly interesting if it is possible to show that the formalism in which it is

---

*Any views expressed in this paper are those of the author. They should not be interpreted as reflecting the views of The RAND Corporation or the official opinion or policy of any of its governmental or private research sponsors. Papers are reproduced by The RAND Corporation as a courtesy to members of its staff.

This paper was prepared for presentation at the National Science Foundation Seminar on the Construction of Complex Grammars, Harvard University, Cambridge Massachusetts, June 1970.

stated would be inadequate to characterize other sets of pairs which do not, however, constitute human languages. To the extent that the formalism is inapplicable except to languages which either are or might plausibly be ordinary human languages, it constitutes a statement about the nature of human language as a whole.

Since language seems to be, before all else, a system in which meanings are encoded in sound, and since languages differ from one another precisely in the particular meaning—sound pairs that they contain, then the study of linguistic competence is the study of what is most essential about language and is logically prior to any other study. However, there are a great many questions that can legitimately be asked about language but which go beyond, in one way or another, the search for an abstract characterization of meaning—sound pairs. It is usual to group these questions under the general heading of "performance" despite the fact that they do not exhibit any other interesting common properties. For example, questions about the sequences of sounds people actually utter as opposed to those that an ideal speaker of the language might be supposed to utter are questions about performance. If a method of characterizing languages is sought which makes it simpler to find the meanings that correspond to a particular sequence of sounds or the sequences of sounds that correspond to a particular meaning, then what is afoot is an investigation of linguistic performance.

Slips of the tongue, the sets of sentences that people find difficult to understand, and segments of text which embrace more than a single sentence all belong to the study of performance. In this paper I shall discuss in very broad outline a particular kind of investigation of certain aspects of linguistic performance that interests me. I shall try to show what I think such an investigation has gained and might still gain from studies of linguistic competence and some implications that I think this work has for linguistics as a whole.

In an article in Mind in 1950, Alan Turing suggested as a worthy scientific goal the construction of a machine with which it would be possible to converse, presumably by telephone or telegraph, and which would be capable of masquerading as a human being. In other words, there would be no way that a person could discover, through simple conversation with the machine, that it was not a real human being. The question of whether such a machine could be constructed in principle is one with far reaching implications, but it is a question which will not be answered soon. In any case, the attempt to build such a machine can be made interesting irrespective of whether the ultimate goal is ever reached. The capabilities that a machine would have to have in order to pass itself off as human are, of course, enormously varied. The machine must clearly be able to understand ordinary sentences and to formulate new

ones to represent the meanings it has to convey.

A first step towards the design of the kind of machine Turing envisaged would be a somewhat different machine whose capabilities, while considerably more modest, would still require a measure of intellectual heroism on the part of the designers. The kind of machine I have in mind is one that would be capable of doing the so-called comprehension exercises that were once popular in the teaching of foreign languages. The student is caused to read a text in the foreign language, and is then required to answer questions, usually also in the language, about what he has read. If it were the nursery rhyme, "Jack and Jill went up the hill", then a possible question might be, "On the injuries to which of the protagonists in this story are we given the most detailed information?" But this is considerably more subtle than the questions usually asked and, in any case, it contains the word "protagonist" which a student at this level may perhaps not be expected to know. More plausible would be questions like, "What did Jack break?", "Who came tumbling after?", and possibly even "Who went to fetch what?". In this exercise the student is supposed to show that he understands that part of the language that has been taught so far, and, to the extent that he answers the questions correctly, it seems to me reasonable to claim that that is exactly what he does show. If the students were given the text of

"Ten Little Indians", then one could imagine asking something like "What proportion of the original number of Indians remained alive after a big bear hugged one?". If a question of this kind brought an outrageous reply, or no reply at all, then the questioner might be led to doubt that he was in conversation with a real human being. On the other hand, he might equally well be led to doubt that he was in conversation with a machine because one thing he does expect machines to be able to do is simple arithmetic calculations. In any case, the ability to answer this kind of question correctly certainly turns on something other than purely linguistic abilities and it therefore seems reasonable to defer treatment of it until later.

Clearly, a machine that understands what people say must be capable of analyzing the grammatical structure of the utterances it receives, representing the information that they contain in some kind of canonical form, of storing expressions in this canonical form in such a way that they can easily be retrieved, of locating those expressions that are relevant to a given question, and of constructing answers that accord with the rules of English grammar. At least in the initial stages of work on a project of this kind, there is little to be gained from drawing a distinction between utterances and sentences. If, in its initial incarnation, the machine is prepared to

converse only with people who express themselves in complete sentences, this will not be accounted a major defect. From now on, I shall therefore talk of the input to the machine as though it consisted entirely of well—behaved sentences, and shall claim, nevertheless, to be engaged in a thorough—going investigation of linguistic performance. To say that we are studying performance is not to say that we abandon the right to make the kinds of idealization of our subject matter that science thrives on.

The reason for wanting to reduce the great variety of sentences that the machine must expect to canonical forms is clear. Whatever questions can be answered wholly or partly on the basis of the sentence, "John gave the book to Mary" can equally well be answered on the basis of sentences like "Mary was given the book by John", "The book was given to Mary by John" and so on. In other words, the system will be more efficient if families of sentences which are related to one another systematically but which do not differ in meaning are made to fall together at the earliest possible stage in the analysis. All the other standard arguments for distinguishing deep from surface structure also go through in the design of this kind of machine. It is clearly desirable, for example, to find a form for sentences in which the logical relationships between words and phrases are represented in a perspicuous manner. At present, it seems that the form of deep

structure that will suit our purposes best will be similar to that suggested by Fillmore, but nothing of what I am saying here turns on this. Considerations that we entertain tomorrow may lead us to believe that the kind of deep structures suggested by the so—called generative semanticists will suit our purposes better.

If we are in fact able to make the machine find the deep structures that correspond to the surface strings it receives, then the efficiency of all subsequent operations will be immeasurably increased, perhaps enough to make practically feasible a project that otherwise would have been totally unthinkable. But efficiency is, in the last analysis, the only thing we can expect to gain. The ideal situation in which all strings with the same meaning would be reduced to exactly one canonical form is something we know to be unachievable. Alan Turing's most famous contribution to mathematics is a theorem, one of whose implications is that we cannot hope to construct an algorithm that will be able to tell us whether any pair of sentences do or do not have the same meaning. Strictly speaking, the theorem is to the effect that the notion of the meaning of a sentence is not entirely coherent, but this is a point to which we shall return.

In a paper entitled "From Semantics to Syntax"* I described a form of rule that would map P—markers onto expressions in a language akin to the predicate calculus.

---

*P—3746, The RAND Corporation, December 1967.

The elementary constants in that language were to be chosen from a closed set of semantic primitives. At the beginning of the paper, I argued that the semantic component of any grammar must furnish readings stated in terms of a universal alphabet of semantic primitives. I take it that, though versions of generative grammar differ in the form of semantic representations they provide, the view that they should be stated in terms of such a universal set of primitives is the one most commonly held. Despite the popularity of this view, I found it necessary in my earlier paper to launch an attack with characteristic linguistic intemperance, on the only other view I know which Lyons and others call <u>structural semantics</u>. I say that, "While this view may have some short-term philosophical appeal, it is almost totally barren from the standpoint of scientific linguistics". Well, having undertaken to build a machine that can do comprehension exercises, I now find myself forced to retreat from my previous position and to espouse structural semantics.

Structural semantics can be characterized briefly as follows. Because a sentence is meaningful, we are led to suppose that there must be something that is its meaning. However, this is not a valid inference. Furthermore, the attempt to isolate meanings from the sentences that normally embody them invariably leads to unsatisfactory results. We can not exhibit the meaning of any sentence otherwise than by writing one or more other sentences in the same or

a different language. It is difficult to see how this gets us any closer to meanings in the pure state. But the structural semanticist claims that there is no isolable entity which can reasonably be called the meaning of a word or phrase, but words and sentences do contract relationships with other words and sentences and, if anything is the true meaning, it is the total set of these relations. If I can say "I saw a tulip", then I can, with equal truth, say "I saw a flower". According to the traditional view we should say that the second sentence can be inferred from the first by virtue of the meanings of the words. According to the structural semanticist, it is exactly the other way about: the fact that the second sentence follows from the first is itself part of the meanings of the words "tulip" and "flower".

Semantic primitives are an intuitively appealing notion and several arguments can be adduced in their favor. If there is any part of linguistic theory in which it seems reasonable to look for universals, it is surely in semantics. If it is possible to translate sentences from one language into another, it is presumably because there is something that is invariant as between a sentence and its translation—what we should normally call its meaning. Semantic primitives also seem to occupy a comfortable place in the area where semantics and syntax overlap. Whilst they belong to semantics, they can, at the same time, be embedded in

contextual features, transformational rules, and the like.

On the other hand, the attempt to actually identify semantic primitives never leads to satisfying results. Arguments can always be found for decomposing proposed primitives into still more elementary units, but the process is governed by a law of diminishing returns so that, with each new round of decomposition, the cash value of the new primitives becomes less and less. Eventually, the process turns full circle and we have as many primitives in the end as we had lexical items in the beginning. Like Katz and Fodor, we find ourselves forced to establish a set of distinguishers along side our semantic markers to capture aspects of meaning which are crucial but totally idiosyncratic.

Lyons points out in his book on structural semantics that the languages of the world contain many sets of words whose meanings come entirely from the relations they contract with other members of the set. The color words are an obvious example. It is well known that different cultures and different languages divide the spectrum into different numbers of units and at different places. Furthermore, for obvious reasons, the boundaries between one color and another can never be sharply defined. The only thing that is absolutely fixed is the sequence in which the colors are encountered as we move from one end of the spectrum to the other. In terms of semantic primitives,

there is very little that can be said about a word like "red" beyond the fact that it stands for a color and that the color in question is red.

Having claimed that meaning relations are logically prior to meaning, rather than the reverse, Lyons goes on to list a number of meaning relations which he considers to be fundamental. These contain, for example, the relation of incompatibility which holds between words like "black" and "white" which cannot both be true of the same thing; and antymony which holds between words like "long" and "short" because to say that something is not long implies that it is short and vice versa; the relation of hyponomy which holds between "mammal" and "dog" because to say that Fido is a dog is to say, by implication, that he is a mammal. Other fundamental relations are those of implication and equivalence. These terms are to be understood as meaning something similar to what logicians intend by them, but there are also crucial differences. From "All men are mortal" and "Socrates is a man" we may infer that Socrates is mortal both in logic and in natural semantics. But the sentence "Two and two equal five" implies that "The moon is made of green cheese" in logic but not in semantics. In other words, by "implication" we do not mean "material implication"; just what we do mean is considerably less clear than we should like it to be.

In my view, relations such as incompatibility, antymony, and hyponomy can all be shown to be parasitic on the notion of implication, and they can therefore be dispensed with. Suppose that our comprehending machine has been given some information about animals and that it has recorded this by storing the canonical deep structure of each of the sentences that it heard. If the machine knows that a dog is a mammal, it is presumably because it was presented with a sentence to that effect, or with a number of sentences from which the fact can be inferred. Suppose that the machine also knows that Fido is a dog and that Fido is in the livingroom. If we now ask the machine whether there are any mammals in the livingroom, it should be able to answer "yes". We could say that it does this in part on the basis of a relation of hyponomy between "mammal" and "dog" and possibly also between "dog" and "Fido" and that it makes the correct inference on the basis of these two relations together with the sentence "Fido is in the livingroom." But, since the machine must be capable of constructing inference schemes based on sentences, there seems to be no reason why it should also take special account of such relations as incompatibility, antymony, and hyponomy. Relationships of these kinds can presumably only be established on the basis of sentences that the machine receives, and it is not clear that anything is to be gained by giving them special status.

What we are proposing, then, is a machine, and, by implication, the theory to underlie that machine, in which the distinction between a speaker's semantic competence and his encyclopedic knowledge of the world is totally obliterated. Intuitively, the appeal of such a scheme is considerable for no amount of introspection can reveal a distinction between semantic properties and facts about the world. The scheme is attractive to someone who wants to build the kind of machine we have had in mind because it reduces the number of parts that the machine must have. To do its job, the machine must clearly be able to remember facts and to make inferences so that it will be able to answer questions, the answer to which is not contained in any single sentence that it has stored.

The scheme is disturbing because it rests heavily on a notion of implication different from material implication and with properties which are, at best, unknown and, at worst, contradictory. But, while we may be disturbed to find that there is much illogic in the logic that underlies ordinary speech, we should surely not be altogether surprised. Indeed, much of what we know about logic was suggested to us by the very illogic of everyday language. There is nothing semantically amiss about a phrase like "the set of all sets" or a word like "self-predicability". But, as we know, the attempt to give them strict logical interpretations leads to contradictions which have,

furthermore, served to illuminate the subject as a whole. The kind of logic we need for semantics is, therefore, unsatisfactory from a mathematicians point of view on at least two counts: first, it does not make use of material implication, and second, it contains no hierarchy of types.

It is worth noting that the notion of synonomy plays no part in the view of semantics I am advocating. This must surely be accounted an advantage because, while there is little difficulty in deciding whether one term is a hyponym of another, or even whether two terms are antonyms, the decision as to whether a pair of terms is synonomous or not is notoriously difficult. In the kind of scheme proposed here, we could say that a pair of terms or a pair of sentences was synonomous if, and only if, they had contracted identical relations with other terms in sentences. It is unlikely that this would ever be the case and, in any case, to verify it would be practically, if not theoretically, impossible.

There may be many people who could countenance the idea of a logic which is substantially illogical and an account of semantics that has nothing to say about synonyms. But if the design of this machine has implications for linguistic theory, they are even more fundamental than these. As I said at the outset, the goal of linguistic theory is to characterize that subset of the Cartesian product of the set of all possible meanings with a set of

all utterable sounds which constitutes a given language. But, if my view of semantics is the correct one, then it no longer makes sense to talk about the set of all possible meanings. In this case, what is left for linguists to do? My answer—and I give it in all possible seriousness—is that linguists should engage in the theoretically revealing, scientifically respectable, practically useful, and altogether exciting enterprise of constructing machines of the kind we have been talking about.

It is possible to take the view that, if there is to be any interest in this enterprise at all, it will come from restrictions that the linguist places on the kinds of mechanisms and components that he allows to be used in building the machine. If no restrictions are imposed, then the machine will have the same theoretical power as a Turing machine, and the claim that it is able to talk and understand will be of relatively little interest. If, on the other hand, the mechanisms that the machine embodies are severely restricted in some way, in particular, if the power of the machine as a whole is considerably less than that of a Turing machine, then the claim that it can speak and understand becomes a much more interesting one, and the theoretical principles underlying the design of the machine become contributions to linguistic theory.

For my own part, I look forward eagerly to a time when we shall need arguments of this metaphysical refinement, but I do not expect that time to come during my life.

It is true, from a purely theoretical point of view, that a machine can work properly without being based on correct principles and without revealing anything of interest about that part of the world that it is being used to model. But, it is also true that any model of human linguistic performance that can, in any reasonable sense, be said to work, is unlikely to be based on wholely falacious principles. The argument is old and simple: if the matter is complicated, then anything that works cannot be all wrong. It is an argument that is used throughout science, and with every good reason. The current view of particle physics is accepted mainly because it works; it is a coherent story about something that it is difficult to write a coherent story about. Now, there is certainly a requirement that entities should not be multiplied beyond necessity, but it is not a difficult requirement to honor. Surely, linguistics cannot long endure if it does not learn to relax a little and to reduce the amount of its resources that it expends on self—conscious brooding over universals to a more reasonable proportion.

To the builder of mechanical models of linguistic processes, the way ahead looks hard, but the outlook is considerably different than from the point of view of theoretical linguistics. Like the linguists of the thirties, we find ourselves less concerned with characterizing languages and separating sentences from non—sentences than

with making some kind of coherent statement about the sentences that do occur. Our grammars have deep structures, but there is no place in them for a base component, if by that we mean a component that is creative and not transductive. Deep structures are whatever the transductive rules yield, given the strings input to the system. If the strings do not make sense, then the machine will probably not be able to make sense of them, and there's and end to it.

The transducer uses not transformations but a highly embellished form of general rewriting rule. This is because the formalism of transformational rules does not allow one to pass from surface strings to underlying structures, but only from P-markers to other P-markers. But, since we have explicitly renounced any claim to represent linguistic universals on the part of the formalism, we have no cause to be self-conscious about this. Finally, since we do not recognize any boundary between semantics and encyclopedic knowledge, we consider that the territory we are licensed to hunt in is unrestricted. There are no linguistic facts and non-linguistic facts. But this is not over-pretentious. At the last trumpet, physics hopes to be ready with an account, in terms of elementary particles, of everything under the sun; we are proposing only to be ready with a sketch of the format that that account is stored in the physicist's mind and some procedures that he may find useful in delivering the report.